

FORM PTO-1390 (REV 11-2000)	U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE	ATTORNEY'S DOCKET NUMBER 36-1489
TRANSMITTAL LETTER TO THE UNITED STATES DESIGNATED/ELECTED OFFICE (DO/EO/US) CONCERNING A FILING UNDER 35 U.S.C. 371		U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.5) 09/913463 Unknown
INTERNATIONAL APPLICATION NO. PCT/GB00/00783	INTERNATIONAL FILING DATE 3 March 2000	PRIORITY DATE CLAIMED 15 March 1999 20 May 1999

TITLE OF INVENTION

RESOURCE SCHEDULING

APPLICANT(S) FOR DO/EO/US

WADDINGTON

Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:

1. ☒ This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
2. ☐ This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
3. ☒ This is an express request to begin national examination procedures (35 U.S.C. 371(f)). The submission must include items (5), (6), (9) and (21) indicated below.
4. ☒ The U.S. has been elected by the expiration of 19 months from the priority date (Article 31).
5. A copy of the International Application as filed (35 U.S.C. 371(c)(2)).
 - a. ☒ is attached hereto (required only if not communicated by the International Bureau).
 - b. ☒ has been communicated by the International Bureau.
 - c. ☐ is not required, as the application was filed in the United States Receiving Office (RO/US).
6. ☐ An English language translation of the International Application as filed (35 U.S.C. 371(c)(2)).
 - a. ☐ is attached hereto.
 - b. ☐ has been previously submitted under 35 U.S.C. 154(d)(4).
7. ☐ Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3))
 - a. ☐ are attached hereto (required only if not communicated by the International Bureau).
 - b. ☐ have been communicated by the International Bureau.
 - c. ☐ have not been made; however, the time limit for making such amendments has **NOT** expired.
 - d. ☐ have not been made and will not be made.
8. ☐ An English language translation of the amendments to the claims under PCT Article 19 (35 U.S.C. 371(c)(3)).
9. ☒ An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
10. ☐ A English language translation of the annexes of the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11 To 20 below concern document(s) or information included:

11. ☐ An Information Disclosure Statement under 37 C.F.R. 1.97 and 1.98.
12. ☒ An assignment document for recording. A separate cover sheet in compliance with 37 C.F.R. 3.28 and 3.31 is included.
13. ☒ A FIRST preliminary amendment.
14. ☐ A SECOND or SUBSEQUENT preliminary amendment.
15. ☐ A substitute specification.
16. ☐ A change of power of attorney and/or address letter.
17. ☐ A computer-readable form of the sequence listing in accordance with PCT Rule 13ter.2 and 35 U.S.C. 1.821-1.825.
18. ☐ A second copy of the published international application under 35 U.S.C. 154(d)(4).
19. ☐ A second copy of the English language translation of the international application under 35 U.S.C. 154(d)(4).
20. ☐ Other items or information.

U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.5) 09/913463 Unknown	INTERNATIONAL APPLICATION NO. PCT/GB00/00783	ATTORNEY'S DOCKET NUMBER 36-1489
---	--	--

21. ☒ The following fees are submitted:

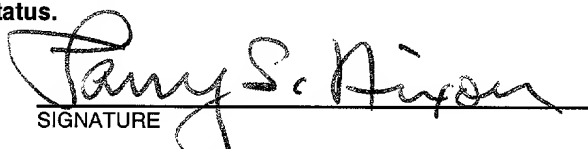
BASIC NATIONAL FEE (37 C.F.R. 1.492(a)(1)-(5):	CALCULATIONS PTO USE ONLY																
-- Neither international preliminary examination fee (37 C.F.R. 1.482) nor international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO and International Search Report not prepared by the EPO or JPO\$1000.00 -- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but International Search Report prepared by the EPO or JPO.....\$860.00 -- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO\$710.00 -- International preliminary examination fee (37 C.F.R. 1.482) paid to USPTO but all claims did not satisfy provisions of PCT Article 33(1)-(4).....\$690.00 -- International preliminary examination fee (37 C.F.R. 1.482) paid to USPTO and all claims satisfied provisions of PCT Article 33(1)-(4).....\$100.00																	
ENTER APPROPRIATE BASIC FEE AMOUNT =	\$	860.00															
Surcharge of \$130.00 for furnishing the oath or declaration later than <input type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 C.F.R. 1.492(e)).	\$	0.00															
<table border="1" style="width:100%; border-collapse: collapse;"> <thead> <tr> <th style="width:20%;">CLAIMS</th> <th style="width:20%;">NUMBER FILED</th> <th style="width:20%;">NUMBER EXTRA</th> <th style="width:20%;">RATE</th> <th style="width:20%;"></th> </tr> </thead> <tbody> <tr> <td>Total Claims</td> <td>16</td> <td>-20 =</td> <td>0</td> <td>X \$18.00</td> </tr> <tr> <td>Independent Claims</td> <td>5</td> <td>-3 =</td> <td>2</td> <td>X \$80.00</td> </tr> </tbody> </table>	CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE		Total Claims	16	-20 =	0	X \$18.00	Independent Claims	5	-3 =	2	X \$80.00	\$	160.00
CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE														
Total Claims	16	-20 =	0	X \$18.00													
Independent Claims	5	-3 =	2	X \$80.00													
MULTIPLE DEPENDENT CLAIMS(S) (if applicable)		\$270.00	\$	0.00													
TOTAL OF ABOVE CALCULATIONS =		\$	1020.00														
<input type="checkbox"/> Applicant claims small entity status. See 37 CFR 1.27. The fees indicated above are reduced by 1/2.		0.00															
SUBTOTAL =		\$	1020.00														
Processing fee of \$130.00, for furnishing the English Translation later than <input type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 C.F.R. 1.492(f)).		0.00															
TOTAL NATIONAL FEE =		\$	1020.00														
Fee for recording the enclosed assignment (37 C.F.R. 1.21(h)). The assignment must be accompanied by an appropriate cover sheet (37 C.F.R. 3.28, 3.31). \$40.00 per property	+	\$	40.00														
Fee for Petition to Revive Unintentionally Abandoned Application (\$1240.00 - Small Entity = \$620.00)		\$	0.00														
TOTAL FEES ENCLOSED =		\$	1060.00														
		Amount to be: refunded	\$														
		Charged	\$														

a. ☒ A check in the amount of \$1060.00 to cover the above fees is enclosed.
 b. ☐ Please charge my Deposit Account No. 14-1140 in the amount of \$_____ to cover the above fees. A duplicate copy of this form is enclosed.
 c. ☒ The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment to Deposit Account No. 14-1140. A duplicate copy of this form is enclosed.
 d. ☒ The entire content of the foreign application(s), referred to in this application is/are hereby incorporated by reference in this application.

NOTE: Where an appropriate time limit under 37 C.F.R. 1.494 or 1.495 has not been met, a petition to revive (37 C.F.R. 1.137(a) or (b)) must be filed and granted to restore the application to pending status.

SEND ALL CORRESPONDENCE TO:

NIXON & VANDERHYE P.C.
 1100 North Glebe Road, 8th Floor
 Arlington, Virginia 22201-4714
 Telephone: (703) 816-4000


 SIGNATURE

Larry S. Nixon
 NAME

25,640
 REGISTRATION NUMBER

August 15, 2001
 Date

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of

WADDINGTON

Atty. Ref.: **36-1489**

Serial No. **Unknown**

Group:

National Phase of: **PCT/GB00/00783**

International Filing Date: **3 March 2000**

Filed: **August 15, 2001**

Examiner:

For: **RESOURCE SCHEDULING**

* * * * *

August 15, 2001

Assistant Commissioner for Patents
Washington, DC 20231

Sir:

PRELIMINARY AMENDMENT

Prior to calculation of the filing fee and in order to place the above identified application in better condition for examination, please amend the claims as follows:

IN THE CLAIMS

Please substitute the following amended claims for corresponding claims previously presented. A copy of the amended claims showing current revisions is attached.

3. (Amended) A method according to claim 1, wherein said resource is a CPU.

4. (Amended) A method according to claim 1, wherein said resource is a mass storage device.

8. (Amended) A method according to claim 1, wherein said first process is a method comprising the steps of:

generating a one-dimensional reservation request pattern; and

merging the reservation request pattern with a one-dimensional CPU access control pattern, representing empty CPU access time slots and reserved CPU access

09/913463

time slots, without substantially disturbing either the reservation request pattern or the reserved CPU access time slots in the reservation request pattern.

11. (Amended) A method according to claim 9, wherein the one-dimensional CPU access control pattern is generated by a method comprising the steps of:

generating a one-dimensional reservation request pattern; and

merging the reservation request pattern with a one-dimensional CPU access control pattern, representing empty CPU access time slots and reserved CPU access time slots, without substantially disturbing either the reservation request pattern or the reserved CPU access time slots in the reservation request pattern.

12. (Amended) A method according to claim 1, wherein the second process performs a method comprising:

generating a one-dimensional CPU access control pattern, each element of which relates to a quantum of CPU access time; and

at the end of a quantum of CPU access time:

granting access to any pending processes having a priority greater than a predetermined level; and then

if the next pattern element is empty then granting access to a pending process meeting a predetermined prioritisation criterion else granting access to a process identified in the pattern element,

wherein pending processes populate queues having different priorities and access is granted to the process identified in the pattern element when there is not a populated process queue having a higher priority than the queue in which said process is present.

WADDINGTON
Serial No. **Unknown**

15. (Amended) A computer configured to operate in accordance with a method according to claim 1.

16. (Amended) A data carrier containing computer code for loading into a computer for the performance of the method of claim 1.

REMARKS

Attached hereto is a marked-up version of the changes made to the claims by the current amendment. The attached page is captioned "**Version with markings to show changes made.**"

The above amendments are made to place the claims in a more traditional format.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By:



Larry S. Nixon

Reg. No. **25,640**

LSN:lmy

1100 North Glebe Road, 8th Floor
Arlington, VA 22201-4714
Telephone: (703) 816-4000
Facsimile: (703) 816-4100

VERSION WITH MARKINGS TO SHOW CHANGES MADE

3. (Amended) A method according to claim 1 [or 2], wherein said resource is a CPU.

4. (Amended) A method according to claim 1 [or 2], wherein said resource is a mass storage device.

8. (Amended) A method according to claim 1, wherein said first process is a method [according to claim 5, 6 or 7] comprising the steps of:

generating a one-dimensional reservation request pattern; and

merging the reservation request pattern with a one-dimensional CPU access control pattern, representing empty CPU access time slots and reserved CPU access time slots, without substantially disturbing either the reservation request pattern or the reserved CPU access time slots in the reservation request pattern.

11. (Amended) A method according to claim 9 [or 10], wherein the one-dimensional CPU access control pattern is generated by a method [according to claim 5, 6 or 7] comprising the steps of:

generating a one-dimensional reservation request pattern; and

merging the reservation request pattern with a one-dimensional CPU access control pattern, representing empty CPU access time slots and reserved CPU access time slots, without substantially disturbing either the reservation request pattern or the reserved CPU access time slots in the reservation request pattern.

12. (Amended) A method according to claim 1, wherein the second process performs a method [according to claim 10 or 11] comprising:

generating a one-dimensional CPU access control pattern, each element of which relates to a quantum of CPU access time; and

at the end of a quantum of CPU access time:

granting access to any pending processes having a priority greater than a predetermined level; and then

if the next pattern element is empty then granting access to a pending process meeting a predetermined prioritisation criterion else granting access to a process identified in the pattern element,

wherein pending processes populate queues having different priorities and access is granted to the process identified in the pattern element when there is not a populated process queue having a higher priority than the queue in which said process is present.

15. (Amended) A computer configured to operate in accordance with a method according to [any preceding claim] claim 1.

16. (Amended) A data carrier containing computer code for loading into a computer for the performance of the method of [any of claims 1 to 14] claim 1.

RESOURCE SCHEDULING

The present invention relates to resource scheduling in a computer.

Enterprise level general-purpose operating systems, both workstation and
5 server, are traditionally designed for efficient resource utilisation. Applications
share a common set of resources, which are initially scheduled according to their
priority and then in a fair weighted manner. This approach results in both an
efficient and fair division of resources to user-level applications. However, certain
applications, in particular multimedia processing applications, are inherently
10 sensitive to resource availability. As a result, the use of priority-based resource
federation policies often results in such applications failing to meet their processing
goals.

US 5,812,844 discloses a method of CPU scheduling, using the Earliest
Deadline First scheduling of threads. However, there is no notion of resource
15 admission in this invention and therefore guarantees are only best-effort and the
scheduling is only single level, based on threads. Scheduling dependent upon
process priority and/or urgency is not included. EP-A-0 798 638 discloses a
further approach to CPU scheduling, in which a number of processes exist as a
group. In turn the group is 'scheduled' by increasing its priority, although the
20 scheduler cannot guarantee that a process within a group will be scheduled.

Conventionally, relatively complex processing is required of the scheduling
process at the stage where a process is granted access to a resource. The present
invention enables a reduction in the complexity of the processing at the access
grant stage by providing for reservation of resources in advance.

25 According to the present invention, there is provided a method of
administering resource utilisation in a computer, the method comprising the steps
of: running a first process to make a reservation for access to a resource in
dependence on a resource requirement communication from an application
process; running a second process to grant requests for access to said resource
30 from said application process in dependence on said reservation; creating a
scheduling means having a method or methods for processing reservation requests
for a plurality of resources and initiating resource specific reservation processing;
creating a reservation means having a method or methods for making reservations
for access to a resource; said application process calling a method of the

scheduling means, said method taking a first resource access requirement definition as a parameter; said method of the scheduling means calling a reservation method of the reservation means to make a reservation for said application process, the reservation method taking a second resource access
5 requirement definition as a parameter; running a resource specific scheduling process to grant access to a resource in dependence on the reservation made by the reservation means; and utilising said resource for the purposes of said application process. The reservation request may be made by an application *per se* or a sub-process or component thereof. Furthermore, access may be granted to
10 any process of an application or for a particular sub-unit of the application for which the reservation request was made.

Making reservations and granting access at application level assists in embodiments of the present invention intended to operate with legacy code. However, where all applications are written for an operating system employing a
15 scheduling method according to the present invention, component or thread level reservation and access grant is preferred.

Preferably, said method of the scheduling means translates the first resource requirement definition into the second resource requirement definition.. The advantage of this is that the resource request can be defined by a programmer
20 in a hardware independent manner and the scheduling object for a particular platform can translate that programmer-produced resource request in a potentially non-linear manner on the basis of the properties of the particular platform.

A method according to the present invention can be applied to the allocation of CPU time, access to mass storage devices, e.g. hard disk drives, optical disk drives and the like, and memory management.
25

Preferably in the case of mass storage access, said first process comprises allocating one or more "lottery ticket" numbers to said application or a process thereof in dependence on the second resource access requirement definition and storing requests for access to a mass storage device from application processes;
30 generating substantially randomly a lottery ticket number; and if no application process has been allocated said lottery ticket number then passing on to a mass storage device driver process the stored request for access from an application process selected on the basis of a predetermined prioritisation criterion else

passing on to a mass storage device driver process a stored request for access from an application process to which said lottery ticket number was allocated.

The majority of multi-tasking operating systems, including Windows NT, Unix and Linux, use a scheduling algorithm known as Round-Robin. At the end of
5 a time slot, or when a thread has revoked its time slice, the dispatcher takes off the first thread which is waiting on the highest priority queue (some queues will be empty). Additional techniques are used to prevent starvation by dynamically increasing the priority of long-waiting threads.

According to the present invention, there is also provided a method of
10 scheduling access to a CPU which may be implemented in a method of administering resource utilisation in a computer according to the present invention, the method comprising the steps of: generating a one-dimensional reservation request pattern; merging the reservation request pattern with a one-dimensional CPU access control pattern, representing empty CPU access time slots and
15 reserved CPU access time slots, without substantially disturbing either the reservation request pattern or the reserved CPU access time slots in the reservation request pattern.

The reservation request pattern may be shorter than the CPU access control pattern, in which case the reservation request pattern is effectively extended by
20 repetition in the merging process. Preferably, said merging step comprises relocating a non-empty time slot element of the reservation request pattern or the CPU access control pattern such that the patterns can be merged without any reserved CPU access time slot elements being deleted or overwritten. More preferably, the relocated non-empty time slot element is relocated by an amount
25 defined in said time slot element. Both forwards and backwards shifts or shift limits may be included in each element.

According to the present invention, there is provided a method of granting access to a CPU which may be implemented in a method of administering resource utilisation in a computer according to the present invention, the method
30 comprising:

generating a one-dimensional CPU access control pattern, each element of which relates to a quantum of CPU access time; and at the end of a quantum of CPU access time:

granting access to any pending processes having a priority greater than a predetermined level; and then if the next pattern element is empty then granting access to a pending process meeting a predetermined prioritisation (e.g. Round-Robin) criterion else granting access to a process identified in the pattern element.

5 Preferably, an entry for the process in any of a plurality of different priority queues will be searched for and access will be granted in respect of the entry for the process having the highest priority. Alternatively, access is granted to the process identified in the pattern element when there is not a populated process queue having a higher priority than the queue in which said process is present.
10 Preferably, the a one-dimensional CPU access control pattern is generated by a method of scheduling access to a CPU according to the present invention.

Embodiments of the present invention will now be described, by way of example, with reference to the accompanying drawings, in which:-

Figure 1 is a block diagram of a general purpose computer;

15 Figure 2 is a block diagram of software components embodied in a computer employing resource scheduling according to the present invention;

Figure 3 illustrates part of a CPU time reservation method according to the present invention;

Figure 4 is a flow chart illustrating the operation of the CPU
20 secondary scheduler of Figure 2; and

Figure 5 is a flow chart illustrating an alternative manner of operation of the CPU scheduler of Figure 2.

Referring to Figure 1, a general purpose computer comprises a CPU 1,
25 RAM 2, ROM 3, a video interface card 4, a hard disk drive controller 5, a keyboard interface 6 and a mouse interface 7 all interconnected by a combined data and address bus 8. A video monitor 9 is connected to the output of the video interface card 4. A hard disk drive 10 is connected to the hard disk drive controller 5. A keyboard 11 is connected to the input of the keyboard interface 6 and a mouse 12 is connected to the input of the
30 mouse interface 7.

The general purpose computer is operating under the control of a multi-tasking operating system. The operating system may be a known

operating system such as Windows NT, Unix or Linux but with the resource scheduling modified to operate as described below.

Referring to Figure 2, in order to provide resource sharing the operating system comprises a dispatcher component 11, a primary scheduler component 12, a CPU reservation component 13, a hard disk drive reservation component 14, a memory reservation component 15, a CPU secondary scheduler 16, a hard disk drive secondary scheduler 17 and a memory balance manager 19.

The dispatcher component 11 maintains queues of threads awaiting servicing by the CPU 1. A thread may have a priority value between 0 and 31, 31 being the highest priority, and the dispatcher component 11 maintains a separate queue for each priority. Thus, all the threads with priority 0 are in one queue, all the threads with priority 1 are in another queue and so on. Within each queue, new threads are added to the tail of the queue which operates on a FIFO principle.

The CPU secondary scheduler 16 is responsible for granting access to the CPU 1 to threads in a manner that will be described in more detail below. The hard disk drive secondary scheduler 17 runs a lottery to determine access to the hard disk drive 10.

The primary scheduler 12 translates high-level abstract resource reservation requests, made by or on behalf of application components 22, into a form suitable for use by the reservation components 13, 14, 15. The reservation components 13, 14, 15 attempt to reserve the requested resources.

A guaranteed service application 21 for use with the present embodiment, consists of a plurality of a components 22 that are instantiated as necessary. Each component 22 comprises at least one thread 23. The constructor 24 of each component 22 includes a procedure which calls CPU, hard disk drive and memory reservation methods of the primary scheduler 12. The component 22 indicates its level of need for a resource by passing the application's process ID (which is unique to the application 21), an indication of the amount of resource required as parameters when calling the reservation methods and optionally a duration for the reservation. Alternatively, the application 21 itself or a management entity, can make a

reservation on behalf of the guaranteed service application 21. A best effort service application 25 consists of a plurality of a components 26 that are instantiated as necessary. The best effort service application's components 26 do not make resource reservations or have reservations made for them.

5 Considering now the case of a request for CPU time, when the CPU reservation method of the primary scheduler 12 is invoked by an application component's constructor 24, the CPU reservation method maps a percentage of resource time, received as a reservation parameter, into a component "signature". The component signature comprises a one-dimensional array of
10 up to 256 elements, each element of which is a record comprising fields for the application's process ID for the component 22, a reservation ID, the forwards flexibility and the backwards flexibility. Initially, the fields of the elements of the array are empty. The fields of selected elements of the array are then filled using the parameters of the call made by the component
15 22. For instance, if the percentage value in the call is 10%, every tenth element of the array is filled and if the percentage value us 20%, every fifth element of the array is filled in. This mapping generally aims to achieve the finest grained reservation signature possible.

Once the component signature has been completed, the CPU
20 reservation component 13 is instantiated on demand by the primary scheduler 12. When the CPU reservation component 13 is instantiated, the primary scheduler 12 forwards the component signature, application and reservation Ids and any duration to the CPU reservation component 13 via a method invocation.

25 The CPU reservation component 13 then tries to merge the component signature into a copy 31 of a master signature 32. The component signature is replicated to fill the full lengths of the signature arrays 31, 32 used by the CPU reservation component 13 and the CPU secondary scheduler 16. For example a reservation signature of 25% maps
30 onto a 1 in 4 slot signature which is repeated 64 times. If any elements of the component signature clashes with an existing reservations (see Figure 3), the CPU reservation component 13 tries moving the clashing elements of the component signature according to their forwards and backwards flexibility values and moving the clashing elements of the copy 31 of the

master signature 32, according to their backwards and forwards flexibilities, until the component signature fits into the copy 31 of the master signature 32. If this cannot be achieved an exception is raised and communicated back to the component 22 via the primary scheduler 12. If merging is
5 achieved, the master signature 32 is updated according to the modified copy 31 and this information is communicated back to the primary scheduler 12. The primary scheduler 12 then returns to reservation ID to the calling component 22.

The CPU reservation component 13 always maintains at least a
10 predetermined minimum number of elements of the master signature 32 free for the set of best effort applications 25. Each element of the master signature 32 corresponds to one CPU access time slice, e.g. 20ms. The CPU reservation component 13 automatically deletes reservations, requested with a duration parameter, that have time-expired.

15 The CPU secondary scheduler 16 uses the master signature 32 cyclically to allocate CPU time to components 22. The use of allocated CPU time by the threads 23 within a component 22 is the responsibility of the application programmer.

Priority based scheduling is used to schedule threads 23 which are very
20 high priority threads (above priority level 15) that are essential to the integrity of the system. This means that threads 23 associated with time-critical operating system tasks, do not have to undergo the process of reservation and admission, as described above. The master signature 32 is used specifically for priority levels 0-15. The reason for this approach is to make possible support for legacy operating
25 system code and services that are essential to the stability of the system.

The CPU secondary scheduler 16, uses the standard Round-Robin scheduling algorithm for threads with a priority greater than 15. Generally, these threads will not require a complete time-slice and are therefore negligible in terms of CPU access time.

30 Referring to Figure 4, the CPU secondary scheduler 16 repeatedly performs the following process. The CPU secondary scheduler 16 determines whether a new thread is pending in the dispatcher 11 (step s1). If a new thread is pending, the CPU secondary scheduler 16 determines whether a thread is currently running (step s2). If a thread is currently running, it determines whether the new pending

thread has a higher priority than the running thread (in this context a guaranteed service thread has "priority" over a best-effort service thread with the same priority) (step s3). If the new thread's priority is higher, the CPU secondary scheduler 16 pre-empts the running thread, dispatches the new thread and places
 5 the pre-empted thread at the head of its priority level queue in the dispatcher 11 (step s4).

If, at step s1, a new thread is not ready or, at step s2, no thread is running, the CPU secondary scheduler 16 determines whether the current time slice has expired (step s5). If the current time-slice has not expired then the process returns
 10 to step s1 else dispatches the thread (step s7) at the head of the dispatcher queue for the highest priority greater than 15 (step s6).

If, at step s6, no threads with priorities greater than 15 are pending, the CPU secondary scheduler 16 inspects the next element of the master signature (step s8). If the element is empty, the CPU secondary scheduler 16 dispatches
 15 the thread from the front of the highest priority queue that is not empty (step s9). If the element is not empty, the CPU secondary scheduler 16 looks for a thread belonging to the application, identified in the master signature element, in the highest priority populated queue in the dispatcher 11 (step s10). If one is found then it is dispatched (step s11) else the thread from the front of the queue is
 20 dispatched (step s12) and the CPU secondary scheduler 16 attempts to move the reservation (step s13). This dynamic slot shift can only be made to the extent of the next slot reserved by the same guaranteed service application 21. The reservation is returned to its correct position for the next cycle through the master signature and if the master signature is updated by the CPU reservation
 25 component 13.

It should be noted that whenever a guaranteed service thread cannot be serviced at steps s8 and s10, the best-effort service application 25 has the possibility of having its threads 26 dispatched.

When a component 22 of the guaranteed service application 21 is
 30 destroyed, its destructor 27 it may invoke a reservation cancel method of the primary scheduler 12, passing the reservation ID as a parameter. The primary scheduler 12 then invokes a cancellation method of the CPU reservation component 13 which removes the calling component's slot

20250313 15:04:50

reservations from the copy 31 of the master signature 32 and then updates the master signature 32.

Considering now the case of disk access, an IRP (I/O Request Packet) is an encapsulated request, which defines the processing requirements, including the type of operation to be carried out (read or write), the number of bytes to
5 manipulate, the synchrony of the request etc..

The constructor 24 of a component 22 includes a procedure for registering the guaranteed service application 21 for disk access. This process calls a disk access reservation method of the primary scheduler 12 with a demand level
10 indicator, usually a percentile, as a parameter. The primary scheduler 12 then calls a disk access reservation method of the hard disk drive reservation component 14 with the demand level indicator as a parameter. The hard disk drive reservation component 14 then allocates a number of "lottery tickets" to the application 21 containing the component 22 in dependence on the demand level indicator. The
15 number of lottery tickets issued may be increased as the demand level indicator is increased. The component 22 may request an abstract reservation in terms of a class identifier which the primary scheduler 12 can use to reference a persistent lookup table mapping the appropriate demand level indicator.

When a thread 23, belonging to a guaranteed service application 21,
20 requires disk access, it sends an IRP to an operating systems I/O manager 20, which is then forwarded to the hard disk secondary scheduler 17 via the operating system's file system component 30. The operating system's file system component 30 increases the granularity of the access request, e.g. from a request for one large block a data to many small blocks of data. When the IRP reaches the
25 hard disk secondary scheduler 17, it is place in a wait array according to the guaranteed service application's process ID. The hard disk drive secondary scheduler 17 runs a lottery by randomly generating "ticket" numbers. If winning application 21 has an IRP in the wait array, the winning application 21 is granted disk access and the longest waiting IRP for the winning application is passed on to
30 a physical disk device driver process 32. When there is no winning ticket holder or the winning ticket holder does not have an IRP ready for servicing, then an alternative IRP is scheduled via a simple cyclic selection scheme which includes IRPs from the best effort service application 25.

When a component 22 of the guaranteed service application 21 is destroyed, its destructor 27 determines whether it is the only instance of a component 22 of the application 21. If it is the last such instance, it invokes a cancel hard disk access method of the primary scheduler 12 which in turn invokes
5 a cancellation method of the hard disk drive reservation component 14. The hard disk drive reservation component 14 then de-allocates the tickets allocated to the application 21.

Considering now the case of memory access, each application 21, 25 maintains its own unique virtual address space, which is a set of memory addresses available for its threads to use. This address space is much larger than
10 the RAM 2 of the computer. When a component 22, 26 references its application's virtual memory, it does so with a pointer which is used by a virtual memory manager to determine the location in RAM 2 to which the virtual address maps. However, some portions of memory may actually reside on disk 10 or
15 backing store. This means that data or code which is infrequently accessed is held on disk 10 thus retaining RAM 2 for better use. To facilitate this scheme, memory is managed in fixed size blocks known as "pages", which are usually 4Kb or 8Kb. If a component 22 references an address whose page is not in RAM 2, then a page fault occurs. This triggers a process known as "paging" which is the task of
20 locating the faulted page within the page file and then loading the page into RAM 2. If there is insufficient space in RAM 2, then a page must be first swapped out and written to the page file. However, each process usually keeps a minimum number of pages locked into RAM 2 known as the "working set".

All modern operating systems generally protect memory in three forms:-
25 physical hardware disallows any thread from accessing the virtual address space of other components, a distinction is made between the mode of operation, kernel-mode (ring 0) which allows threads access to system code and data, and user-mode (ring 3) which does not, and a page-based protection mechanism wherein
30 each virtual page maintains a set of flags that determine the type of access permitted. However, these mechanisms are aimed at protecting a component's memory resources from unwanted interference by other unauthorised processes in the system. They do not prevent a process from consuming more than its reasonable share of memory.

To aid understanding of the approach of the present invention to the implementation of working set size control, the working set management scheme in Windows NT will be briefly described. When created, each component is assigned two thresholds, a minimum working-set size and maximum working-set size. The minimum defines the smallest number of pages the virtual memory manager attempts to keep locked concurrently in physical memory, whilst the maximum defines an expansion range which can be used if the component is causing a considerable number of page faults. If the working set is too small, then the component incurs a large number of paging operations and thus a substantial overhead is incurred through continuously swapping pages to and from the backing store. Alternatively, if the working set is too large, few page faults occur but the physical memory may be holding code and/or data which is infrequently referenced and thus overall efficiency is reduced. In Windows NT, the Memory Manager (MM) adjusts the working sets once every second, in response to page-in operations or when free memory drops to below a given threshold. If free memory is plentiful, the MM removes infrequently referenced pages only from working sets of components whose current size is above a given minimum, this is known as aggressive trimming. However, if free memory is scarce, the MM can force trimming of pages from any component until it creates an adequate number of pages, even beyond the minimum threshold. Of course, components which have extended working sets are trimmed in preference to those which have the minimum working set size.

In the present embodiment, the constructor 24 of a component 22 includes a procedure for reserving a working set. This procedure calls a memory reservation method of the primary scheduler 12 which, in turn calls a method of the memory reservation component 15. The memory reservation component 15 translates this request into the appropriate expansion of a default minimum working set size. The minimum working set size regulates the number of pages an application 21 can concurrently lock (or pin) into RAM 2. Thus, a component 22 that wishes to reserve a portion of memory, makes the reservation request via the primary scheduler 12 to the memory reservation component 15 which then increases the working set thresholds accordingly. This adjustment is facilitated through operating system calls. It is then the responsibility of the component to pin the pages into RAM 2 as required. The pinning process is arbitrated on a first-

come first-serve basis. If insufficient physical pages are available, the operating system will reject the pin request, therefore the component 22 is expected to pin its reserved pages soon after allocation.

If there are insufficient free physical memory resources, the memory reservation component 15 then attempts to force the processes of best-effort service applications 25 to write pages back to file store until sufficient free physical pages are available. If insufficient space can be made, then the reservation is rejected. As with the other resource modules, a portion of the RAM 2 is reserved for a minimal best-effort service. This is facilitated by maintaining a count of the total number of physical pages being used under the guaranteed service and checking this against a portion of the total physical page capacity. It should also be noted that the 'best-effort' service is strictly a minimal guaranteed service (defined by the system assigned minimum working set size) in addition to the best-effort service through potential expansion of the working set. It is assumed that processes cannot bypass the primary scheduler 12 and directly alter their own working set size. Finally, special consideration is made for pages which are being used for shared memory purposes. If any client of a shared page has undergone reservation, then the page is classified as reserved even though other clients may not have made a reservation.

A second embodiment is the same as the first embodiment described above except for the manner of operation of the CPU secondary scheduler 16. The operation of the CPU scheduler of the second embodiment will now be described.

Referring to Figures 2 and 4, the CPU secondary scheduler 16 repeatedly performs the following process. The CPU secondary scheduler 16 determines whether a new thread is pending in the dispatcher 11 (step s101). If a new thread is pending, the CPU secondary scheduler 16 determines whether a thread is currently running (step s102). If a thread is currently running, it determines whether the new pending thread has a higher priority than the running thread (in this context a guaranteed service thread has "priority" over a best-effort service thread with the same priority) (step s103). If the new thread's priority is higher, the CPU secondary scheduler 16 pre-empts the running thread, dispatches the new thread and places the pre-empted thread at the head of its priority level queue in the dispatcher 11 (step s104).

If, at step s101, a new thread is not ready or, at step s102, no thread is running, the CPU secondary scheduler 16 determines whether the current time slice has expired (step s105). If the current time-slice has not expired then the process returns to step s101 else dispatches the thread (step s107) at the head of the dispatcher queue for the highest priority greater than 15 (step s106).

If, at step s106, no threads with priorities greater than 15 are pending, the CPU secondary scheduler 16 inspects the next element of the master signature (step s108). If the element is empty, the CPU secondary scheduler 16 dispatches the thread from the front of the highest priority queue that is not empty (step s109). If the element is not empty, the CPU secondary scheduler 16 looks for a thread belonging to the application, identified in the master signature element, in a queue in the dispatcher 11 (step s110), searching from the priority 15 queue towards the lowest priority queue and from the head of each queue to the tail thereof. If one is found then it is dispatched (step s111) else the thread from the front of the highest priority populated queue is dispatched (step s112) and the CPU secondary scheduler 16 attempts to move the reservation (step s113). This dynamic slot shift can only be made to the extent of the next slot reserved by the same guaranteed service application 21. The reservation is returned to its correct position for the next cycle through the master signature and if the master signature is updated by the CPU reservation component 13.

It should be noted that whenever a guaranteed service thread cannot be serviced at steps s108 and s110, the best-effort service application 25 has the possibility of having its threads 26 dispatched.

The present invention has been described with reference to applications comprising components which in turn comprise one or more threads. It will be appreciated that the present invention is not limited to this situation and may be implemented in a system in which application programs are not built using object-oriented languages.

Whilst the present invention has been described with reference to CPU, hard disk and memory access, it will be appreciated that it may be applied to the scheduling of access to other resources.

It will be appreciated that the term "lottery ticket" is used figuratively.

CLAIMS

1. A method of administering resource utilisation in a computer, the method comprising the steps of:
- 5 running a first process to make a reservation for access to a resource in dependence on a resource requirement communication from an application process;
- running a second process to grant requests for access to said resource from said application process in dependence on said reservation;
- 10 creating a scheduling means having a method or methods for processing reservation requests for a plurality of resources and initiating resource specific reservation processing;
- creating a reservation means having a method or methods for making reservations for access to a resource;
- 15 said application process calling a method of the scheduling means, said method taking a first resource access requirement definition as a parameter;
- said method of the scheduling means calling a reservation method of the reservation means to make a reservation for said application process, the reservation method taking a second resource access requirement definition as a
- 20 parameter;
- running a resource specific scheduling process to grant access to a resource in dependence on the reservation made by the reservation means; and
- utilising said resource for the purposes of said application process.
- 25 2. A method according to claim 1, wherein said method of the scheduling means translates the first resource requirement definition into the second resource requirement definition.
3. A method according to claim 1 or 2, wherein said resource is a CPU.
- 30 4. A method according to claim 1 or 2, wherein said resource is a mass storage device.

5. A method of scheduling access to a CPU, the method comprising the steps of:

generating a one-dimensional reservation request pattern; and

merging the reservation request pattern with a one-dimensional CPU access control pattern, representing empty CPU access time slots and reserved CPU access time slots, without substantially disturbing either the reservation request pattern or the reserved CPU access time slots in the reservation request pattern.

6. A method according to claim 5, wherein said merging step comprises relocating a non-empty time slot element of the reservation request pattern or the CPU access control pattern such that the patterns can be merged without any reserved CPU access time slot elements being deleted or overwritten.

7. A method according to claim 6, wherein the relocated non-empty time slot element is relocated by an amount defined in said time slot element.

8. A method according to claim 1, wherein said first process is a method according to claim 5, 6 or 7.

9. A method of granting access to a CPU comprising:
generating a one-dimensional CPU access control pattern, each element of which relates to a quantum of CPU access time; and
at the end of a quantum of CPU access time:
granting access to any pending processes having a priority greater than a predetermined level; and then
if the next pattern element is empty then granting access to a pending process meeting a predetermined prioritisation criterion else granting access to a process identified in the pattern element.

10. A method according to claim 9, wherein pending processes populate queues having different priorities and access is granted to the process identified in the pattern element when there is not a populated process queue having a higher priority than the queue in which said process is present.

11. A method according to claim 9 or 10, wherein the one-dimensional CPU access control pattern is generated by a method according to claim 5, 6 or 7.

12. A method according to claim 1, wherein the second process performs a
5 method according to claim 10 or 11.

13. A method of administering resource utilisation in a computer, the method comprising:

running a first process to make a reservation for access to a resource in
10 dependence on a resource requirement communication from an application process; and

running a second process to grant requests for access to said resource from said application process in dependence on said reservation, wherein

said first process comprises allocating one or more resource tokens to said
15 application process in dependence on the second resource access requirement definition and

the second process performs a method comprising the steps of:

- (i) storing requests for access to a mass storage device from application processes;
- 20 (ii) generating substantially randomly a resource token identifier; and
- (iii) if no application process has been allocated said identified resource token then passing on to a mass storage device driver process the stored request for access from an application process selected on the basis of a predetermined prioritisation criterion else passing on to a mass storage device driver process a
25 stored request for access from an application process to which said identified resource token was allocated.

14. A method of administering resource utilisation in a computer, the method comprising:

30 running a first process to make a reservation for access to a resource in dependence on a resource requirement communication from an application process; and

running a second process to grant requests for access to said resource from said application process in dependence on said reservation, wherein

the second process performs a method comprising the steps of:

(ii) using a stochastic process, either selecting an application process a probability determined by the weighting associated with the application process and passing on to a mass storage device driver process the stored request access from the selected application process, or passing on to a mass storage device driver process a stored request for access from an application process selected on the basis of a predetermined prioritisation criterion

15

16. A data carrier containing computer code for loading into a computer for the performance of the method of any of claims 1 to 14.

ABSTRACT

RESOURCE SCHEDULING

5 A method of administering resource utilisation in a computer comprises running a first process (13, 14, 15) to make a reservation for access to a resource in dependence on a resource requirement communication from an application (21) and running a second process (16, 17, 19) to grant requests for access to said resource from said application in dependence on said reservation. A further
10 process (12) provides a common interface between the first process (13, 14, 15) for each resource and the application (21). The further process (12) converts high-level abstract resource requirement definitions into formats applicable to the first process (13, 14, 15) for the resource in question. The processes are preferably implemented at methods of software objects.

15

T05T80" E04E160

09/913463

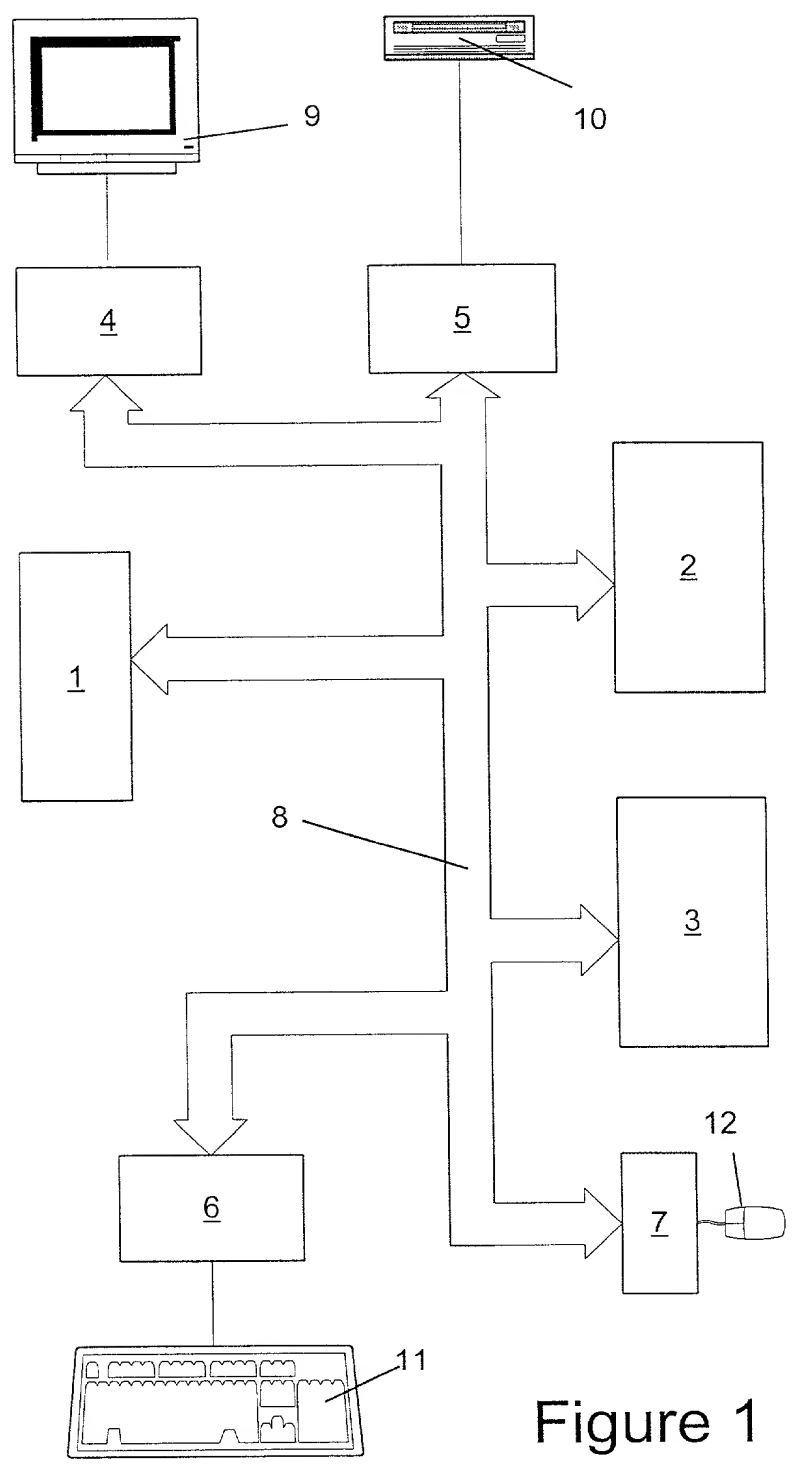


Figure 1




Figure 2




Figure 3

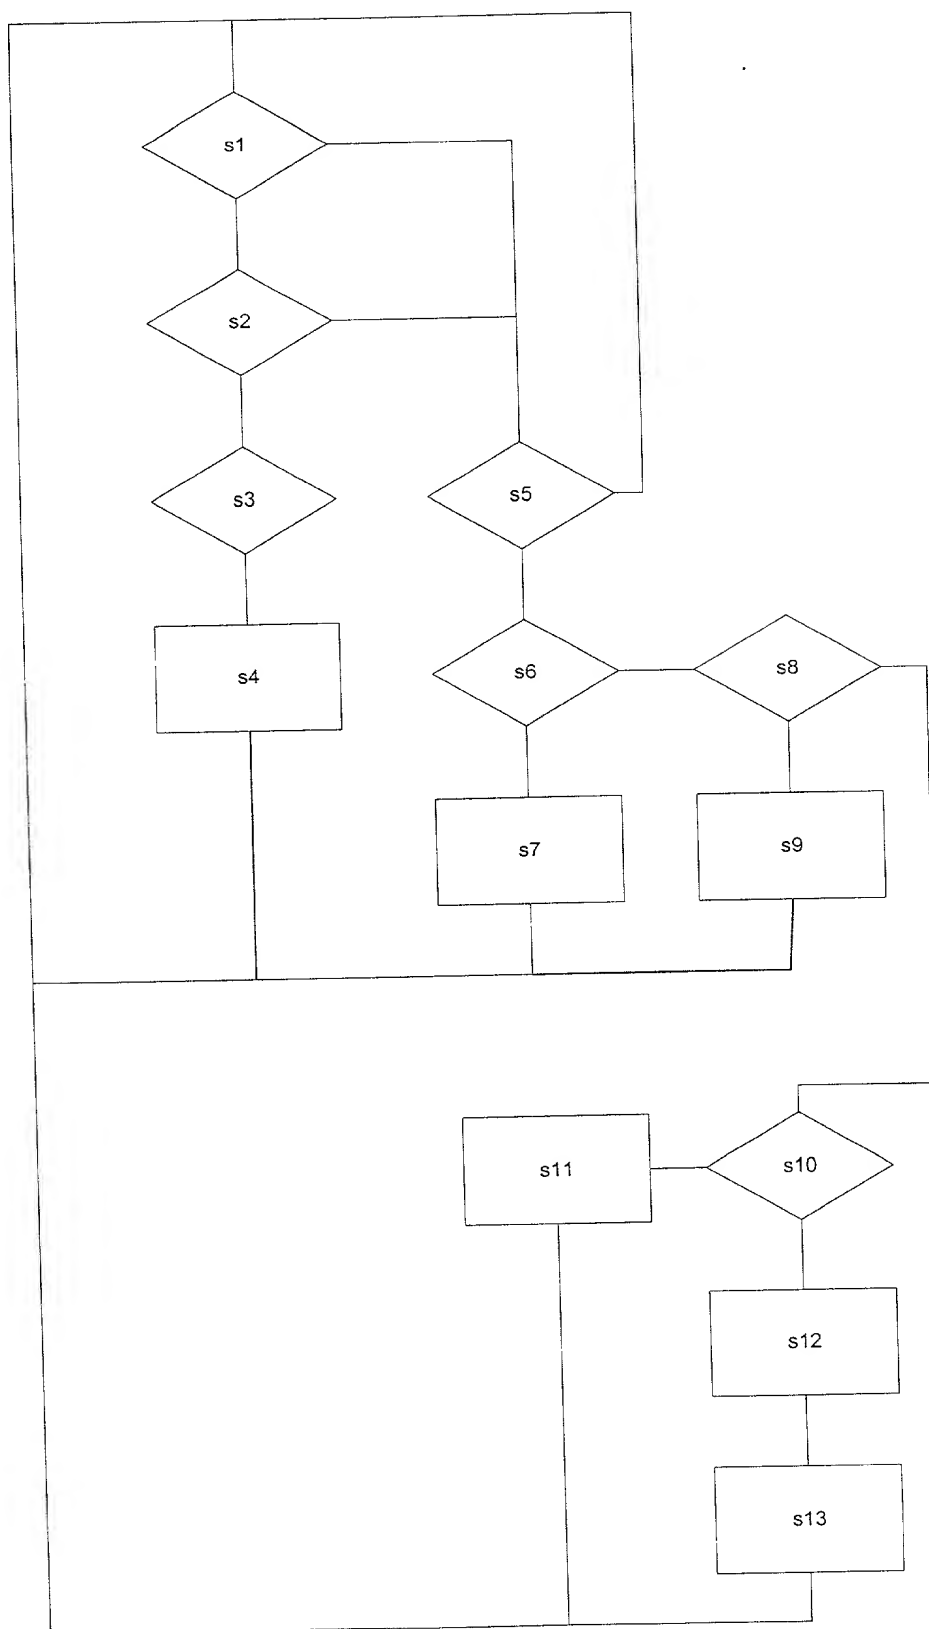


Figure 4

T02130 E042160

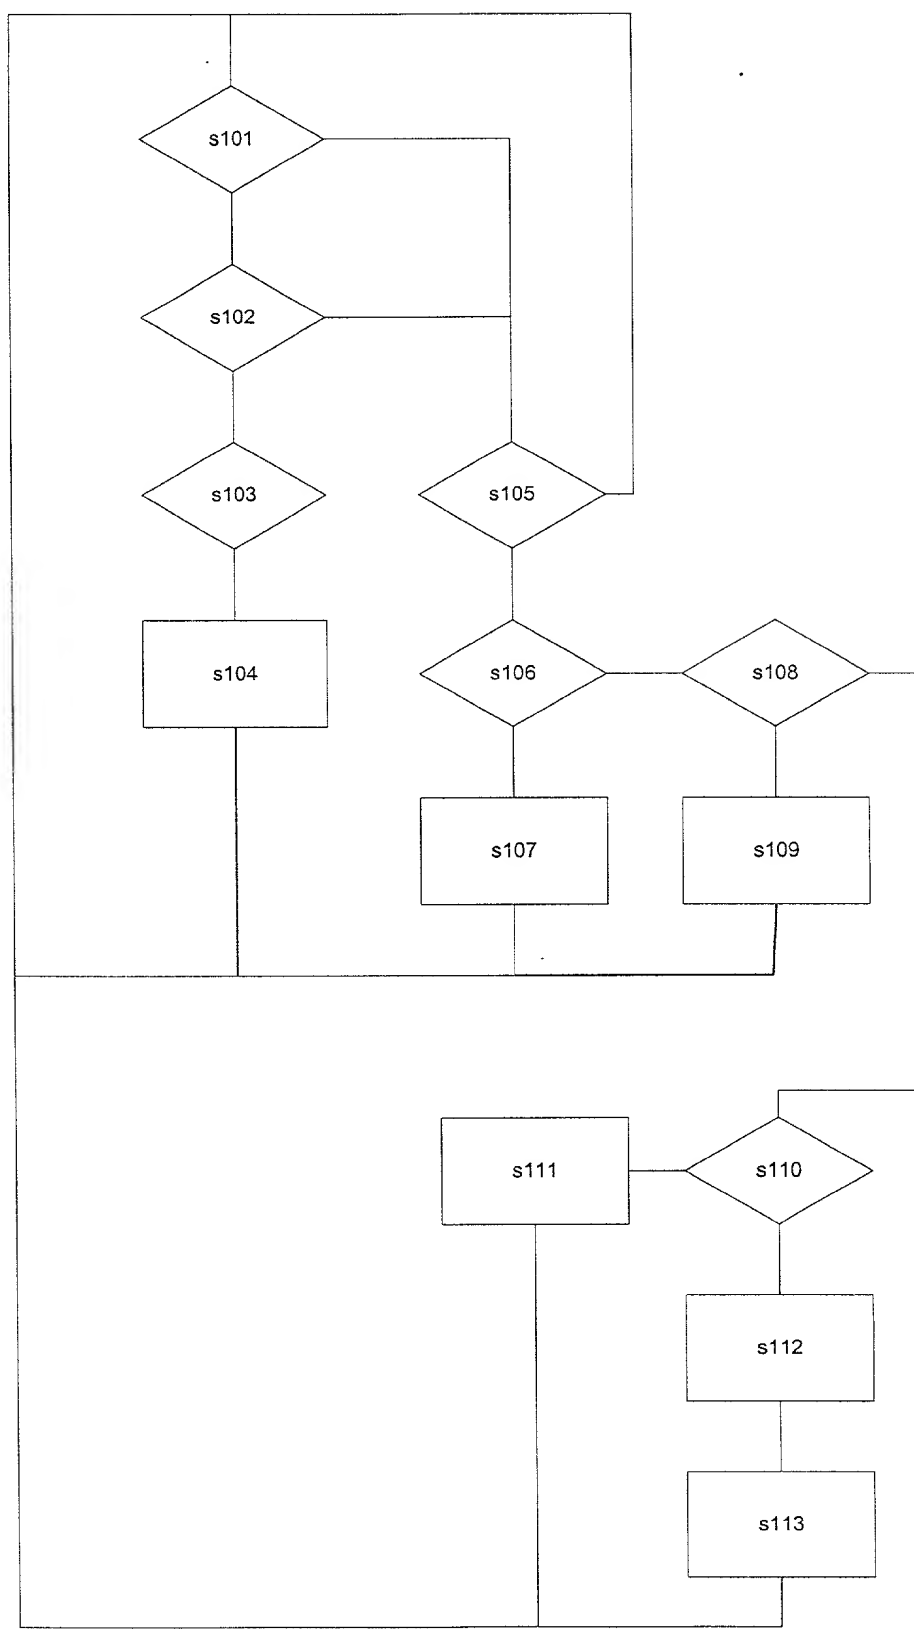


Figure 5

RULE 63 (37 C.F.R. 1.63)
DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

As a below named inventor, I hereby declare that my residence, post office address and citizenship are as stated below next to my name, and I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

RESOURCE SCHEDULING

the specification of which (check applicable box(es)):

- ☐ is attached hereto
☐ was filed on

as U.S. Application Serial No.

(Atty Dkt. No.

-)

☒ was filed as PCT International application No. PCT/GB00/00783 on 03 March 2000

and (if applicable to U.S. or PCT application) was amended on

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 C.F.R. 1.56. I hereby claim foreign priority benefits under 35 U.S.C. 119/365 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed or, if no priority is claimed, before the filing date of this application:

Priority Foreign Application(s):

Application Number	Country	Day/Month/Year Filed
<u>99301952.0</u>	<u>EUROPE</u>	<u>15 MARCH 1999</u>
<u>99303936.1</u>	<u>EUROPE</u>	<u>20 MAY 1999</u>

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

Application Number	Date/Month/Year Filed
--------------------	-----------------------

I hereby claim the benefit under 35 U.S.C. 120/365 of all prior United States and PCT international applications listed above or below and, insofar as the subject matter of each of the claims of this application is not disclosed in such prior applications in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. 1.56 which occurred between the filing date of the prior applications and the national or PCT international filing date of this application:

Prior U.S./PCT Application(s):

Application Serial No.	Day/Month/Year Filed	Status: patented pending, abandoned
<u>PCT/GB00/00783</u>	<u>03 March 2000</u>	<u>PENDING</u>

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon. And on behalf of the owner(s) hereof, I hereby appoint NIXON & VANDERHYE P.C., 1100 North Glebe Rd., 8th Floor, Arlington, VA 22201-4714, telephone number (703) 816-4000 (to whom all communications are to be directed), and the following attorneys thereof (of the same address) individually and collectively owner's/owners' attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith and with the resulting patent: Arthur R. Crawford, 25327; Larry S. Nixon, 25640; Robert A. Vanderhye, 27076; James T. Hosmer, 30184; Robert W. Faris, 31352; Richard G. Besh, 22770; Mark E. Nusbaum, 32348; Michael J. Keenan, 32106; Bryan H. Davidson, 30251; Stanley C. Spooner, 27393; Leonard C. Mitchard, 29009; Duane M. Byers, 33363; Jeffry H. Nelson, 30481; John R. Lastova, 33149; H. Warren Burnam, Jr. 29366; Thomas E. Byrne, 32205; Mary J. Wilson, 32955; J. Scott Davidson, 33489; Alan M. Kagen, 36178; Robert A. Molan, 29834; B. J. Sadoff, 36663; James D. Berquist, 34776; Updeep S. Gill, 37334; Michael J. Shea, 34725; Donald L. Jackson, 41090; Michelle N. Lester, 32331; Frank P. Presta, 19828; Joseph S. Presta, 35329 I also authorize Nixon & Vanderhye to delete any attorney names/numbers no longer with the firm and to act and rely solely on instructions directly communicated from the person, assignee, attorney, firm, or other organization sending instructions to Nixon & Vanderhye on behalf of the owner(s).

1. ☒ Inventor's Signature: [Signature] Date: 13 APRIL 2000
Inventor: DANIEL G WADDINGTON GB
(first) (MI) (last) (citizenship)
Residence: (city) LANCASTER GB (state/country) GREAT BRITAIN
Post Office Address: FLAT 1, 25 ST. GEORGES QUAY, LANCASTER, LANCASHIRE
(Zip Code) LA1 4YR

2. Inventor's Signature: _____ Date: _____
Inventor: _____
(first) (MI) (last) (citizenship)
Residence: (city) _____ (state/country) _____
Post Office Address: _____
(Zip Code) _____

FOR ADDITIONAL INVENTORS, check box ☐ and attach sheet with same information and signature and date for each.